The Harmonic Series

(This article assumes basic knowledge of musical intervals and triadic harmony.)

What is a "sound"? Most people have some concept of invisible waves of air molecules that make their way across a room and tickle little hairs in your ears, which then send a signal to your brain that is interpreted as a "sound."

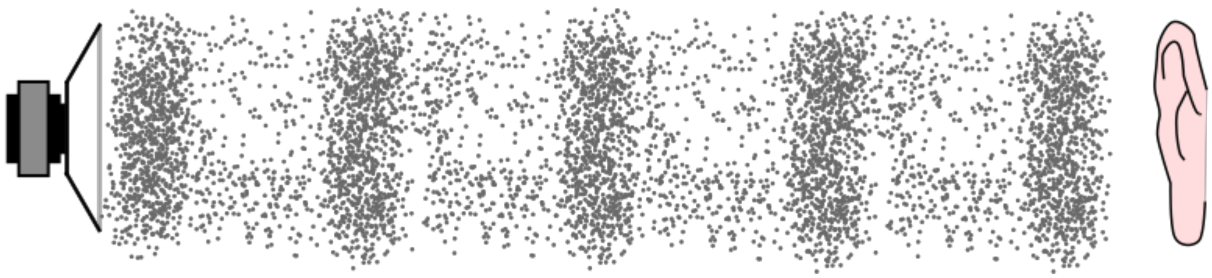

Source: Wikimedia Commons

Sound waves are a pattern in which air molecules are pushed up against each other, temporarily creating an area of increased pressure, and then relieve that pressure by pushing away from each other, creating an area of negative pressure in the space between. This pattern recurs until all the energy has been distributed and equilibrium is reached.

You are probably used to seeing this process represented with a shape like this:

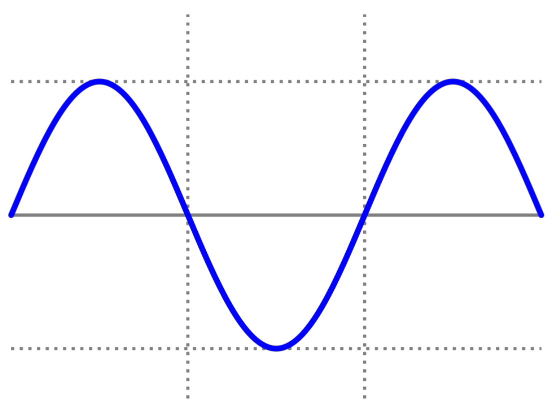

Source: Wikimedia Commons

This is not a representation of what a sound wave would actually look like if we could see it, but a graph in which the x-axis represents time (usually on an imperceptibly small scale) and the y-axis represents pressure. The height of the wave's peaks and valleys is called **amplitude** in scientific terms, and is perceived in our minds as **volume**. Higher-amplitude waves sound louder to us, and lower-amplitude waves sound quieter. The horizontal size of the wave, or the measure of how many times the wave cycles per second, is called **frequency**, and is perceived by our brains as **pitch**. Higher-frequency waves, or waves that cycle through their pattern more times per second, sound higher pitched, and lower-frequency waves sound lower-pitched.

The thing about sound is that it is never just one wave. In fact, you've never heard one single, isolated sound wave in your entire life. The air is full of waves of many different frequencies, loud and soft and constantly changing and interacting with each other. Even something you conceive of as one simple "sound" from a single source (such as clapping your hands) is actually made of many component waves, which are called **partials**.

What makes a sound a musical "note"? If I were to take a mallet and strike the key of a marimba, it would make a very different sound than if I were to strike a non-musical object, such as a table. One is clearly a "note" and the other is clearly a "noise."

Like all sounds, a musical note is not one single sound wave, but a complex combination of waves. However, unlike most sounds, a musical note's component waves have very specific relationships to one another.

The pitch of a musical note is determined by the frequency of its lowest partial, called the **fundamental frequency**. The remaining partials, called **overtones**, have frequencies which are whole-number multiples of the frequencies of the fundamental frequency. All partials in a musical sound, including the fundamental frequency, are called **harmonics**.

For example, if a musical note has a fundamental frequency of 100 Hz (cycling 100 times per second), the second harmonic (or first overtone) would have a frequency of 200 Hz, the third harmonic would have a frequency of 300 Hz, the fourth would have a frequency of 400 Hz, and so on. This can be more universally represented as ratios of the fundamental frequency, called the harmonic series:

1:1 – Fundamental frequency (first harmonic)
2:1 – Second harmonic
3:1 – Third harmonic
4:1 – Fourth harmonic
5:1 – Fifth harmonic
6:1 – Sixth harmonic
etc.

Because harmonics generally get quieter as they get higher, the first six or seven harmonics are usually the ones that are really important.

What does this have to do with music? A lot, actually. First and foremost, these frequency ratios correspond to musical intervals. The most important of these is the 2:1 ratio, or the ratio of the second harmonic to its fundamental frequency. In musical terms, this relationship can be described as a perfect **octave**. Most people, even with no musical training, can inherently recognize that pitches separated by octaves have a special relationship. Just tell anyone to listen while you play various B-flats all around the piano keyboard, and as soon as you play an A instead, they will immediately recognize that it is not like the others.

So, if you have a pitch whose frequency is 100 Hz, the pitch an octave above that would have a frequency of 200 Hz (double.) To achieve yet another octave higher, you would need to produce a pitch with a frequency of 400 Hz (double again.)

As an aside, this illustrates something really strange about the way our brains interpret frequency as pitch: perceived intervals between pitches are not fixed distances, but ratios. Pitches that are 10 Hz and 5 Hz are exactly an octave apart because the ratio between them is 2:1, and pitches that are 17,104 Hz and 8,552 Hz are also exactly an octave apart for the same reason. We hear the "distances" between these two pairs of pitches as equivalent because their ratios are the same, even though their mathematical differences are not even close to the same.

This means that even though the top two keys of a standard piano have about a 235 Hz difference in their frequencies, and the bottom two keys have only about a 2 Hz difference in their frequencies, both of those differences are perceived as an identical interval (1 semitone, or a minor 2nd) because the ratios are the same (roughly 16:15).

This logarithmic scale can be observed in the way the finger positions on string instruments become closer together the higher you move on the fingerboard or fretboard; an octave above the sound of an open string can be achieved by pressing the string at its center, shortening the vibrating segment by half. (Half the string length means double the frequency of vibration, or one octave higher.) The next octave up is three quarters of the way up the string (shortening by half again) and then seven eighths, and so on and so forth, each touch point getting closer and closer together.

If you look at the audible frequency spectrum on an EQ plugin, you can actually see how the view has been adjusted to accommodate the way we perceive ratios rather than simple mathematical differences:

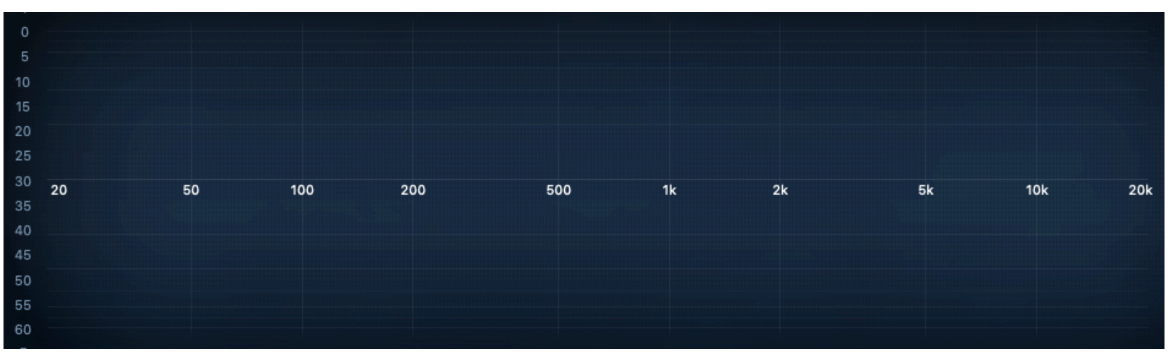

Logic Pro X stock EQ

Notice that the scale of the x-axis is not linear. The distance between 50 Hz and 100 Hz is represented as equivalent to the distance between 10 kHz and 20 kHz. They are the same ratio, and thus sound like the same "distance" (one octave) apart to our ears, a relationship that is represented visually by using a logarithmic—rather than linear—scale.

Some graphical EQ interfaces include vertical lines that give you a sense of how the scale of the x-axis is being represented:

Fabfilter Pro-Q

Let's return to the harmonic series as musical intervals. We've established that a 2:1 frequency ratio is perceived as the interval of an octave. That means that in a musical tone, the second harmonic is an octave higher than the first harmonic, or the fundamental frequency. It also means that the fourth harmonic (4:1) is an octave above the second harmonic (4:2) and two octaves above the fundamental frequency (4:1). Likewise, the eighth harmonic (8:1) will be yet another octave higher.

But there are other ratios. The third harmonic is 3:1 to the fundamental frequency, which means that it is more than an octave (2:1) but less than two octaves (4:1). In fact, it is a perfect 12th, or an octave plus a 5th. (Remember not to count a note twice when adding intervals. In music, 8 + 5 = 12, not 13.) Consider also that if doubling a frequency results in moving up an octave, the opposite is also true. Using this, we can further extrapolate that if 3:1 is the ratio of an octave plus a perfect 5th, we can divide that ratio in half to reduce it by an octave. Therefore, 3:2 (half of 3:1) is the frequency ratio for a perfect 5th.

Here's an example with specific notes: If the fundamental frequency is a pitch called C_2, we know that the second, fourth, and eighth harmonics (2:1, 4:1, 8:1) will be successively higher-pitched notes that are also labeled with the letter C (specifically, C_3, C_4, and C_5) because all notes in the pitch class C are separated by octaves. But because the second harmonic (3:1) is a perfect 12th above the fundamental pitch, it belongs to a different pitch class, i.e. it will be labeled with a different letter. The note a perfect 12th above C (or perfect 5th when reduced to the range of a single octave) is called G.

If you've ever wondered about the foundational importance of the perfect 5th in music (circle of fifths, dominant chords in functional harmony, etc.) this may explain the psychoacoustic origins of those relationships. In fact, the perfect 5th is what determines the number of parts that an octave is divided into. If you move up a 5th (3:2) from a given frequency x, and then move another 5th up from that, and so on until you arrive at something in the same pitch class as x (i.e. $x \bullet 2^y$ where y is a whole number), you will have cycled through a total of twelve pitches, belonging to the twelve pitch classes in Western music. (Well, sort of. More on that later.)

Let's look at the mathematical frequency relationships that we have translated to musical relationships so far:

Partial Name	Ratio to Fundamental	Musical Interval	Octave-Reduced Interval	Example Note
Fundamental Frequency (First Harmonic)	1:1	Unison	n/a	C
Second Harmonic	2:1	8^{ve}	Unison	C
Third Harmonic	3:1	8^{ve} + P5	Perfect 5th	G
Fourth Harmonic	4:1	two 8^{ves}	Unison	C
Fifth Harmonic	5:1			
Sixth Harmonic	6:1			

There are two more harmonics to fill in on this chart, but the sixth harmonic is very simple to find now that we have the third harmonic. Remember, doubling a frequency increases its musical pitch by an octave, so if we know that the third harmonic (3:1) is an octave plus a 5^{th}, we also know that double that ratio (6:1) is <u>two</u> octaves plus a 5^{th}.

Finally, we know that the fifth harmonic (5:1) is somewhere between two octaves (4:1) and two octaves plus a 5^{th} (6:1). In fact, it corresponds to the musical distance of two octaves plus a Major 3^{rd}.

Here is the completed chart:

Partial Name	**Ratio to Fundamental**	**Musical Interval**	**Octave-Reduced Interval**	**Example Note**
Fundamental Frequency (First Harmonic)	1:1	Unison	n/a	C
Second Harmonic	2:1	8^{ve}	Unison	C
Third Harmonic	3:1	8^{ve} + P5	Perfect 5^{th}	G
Fourth Harmonic	4:1	two 8^{ves}	Unison	C
Fifth Harmonic	5:1	two 8^{ves} + M3	Major 3^{rd}	E
Sixth Harmonic	6:1	two 8^{ves} + P5	Perfect 5^{th}	G

Referring to this table, we could say that when you hear a musical note that we called C, you are also (more faintly) hearing several other, higher Cs, as well as a couple of Gs and an E. We can transpose these intervals to any note and say, for example, the musical note A also includes fainter tones of E and C-sharp. You may notice that these harmonics correspond to notes in a major chord rooted in the fundamental frequency. This is a potential scientific explanation for why a major chord sounds "consonant" to our ears when compared to a random group of notes; the frequencies of the notes E and G match up with the overtones of the root note C, so those notes sound "stable" when played together. As with octave relationships, most people, regardless of musical training can identify consonance vs. dissonance, meaning our brains can intuitively identify when groups of pitches have overtones which align with each other. Here is another chart that shows all the partials on a musical staff:

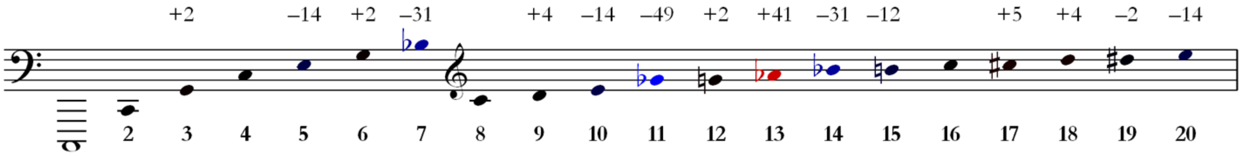

Source: Wikimedia Commons

The harmonics are numbered below each note, with the farthest left being the fundamental frequency. However, the numbers at the top of this chart show a confounding caveat to this system, which is that most of the whole-number ratios result in intervals that are slightly different than the musical intervals we use in modern tuning systems. For example, the third harmonic is commonly referred to as being a perfect 12th above the fundamental, or a perfect 5th when octave-reduced. In fact, it is about 2 cents (or 2% of a semitone) larger than the interval of a perfect 5th as played on a modern piano keyboard. Some of the higher harmonics are much more drastically different from their stated interval, as indicated by blue (much smaller) and red (much larger) noteheads.

If you're extra math-savvy, you may have already noticed a problem back when I mentioned that traveling around the circle of fifths will result in twelve different pitch classes before you return to the one you started with. In truth, traveling by true mathematical fifths (continuously multiplying by 3:2) will never result in a frequency that is exactly in the same pitch class ($x \bullet 2^n$) of the note you started from because no multiple of 3:2 (i.e. 1.5) will ever be evenly divisible by 2.

As illustrated in the notated harmonic series on the previous page, Modern Western tuning, a system called **Equal Temperament**, transgresses on the natural whole-number ratios of the harmonic series by making very small adjustments to many of them. Only the octaves (harmonics 2^n) are left untouched. Without going into a long music theory tangent, I'll just say that this tweaked system was necessary for the development of functional harmony as we know it. (If you have a solid grasp on all of these concepts and want to dive into music theory that feels borderline occult, ask me about the "spiral of true fifths" sometime.)

To review a bit: a musical pitch is not just one frequency, but a fundamental frequency with overtones that follow a whole-number ratio pattern.

Here is the output from a graphical EQ while analyzing a single, sustained C3 by a string section:

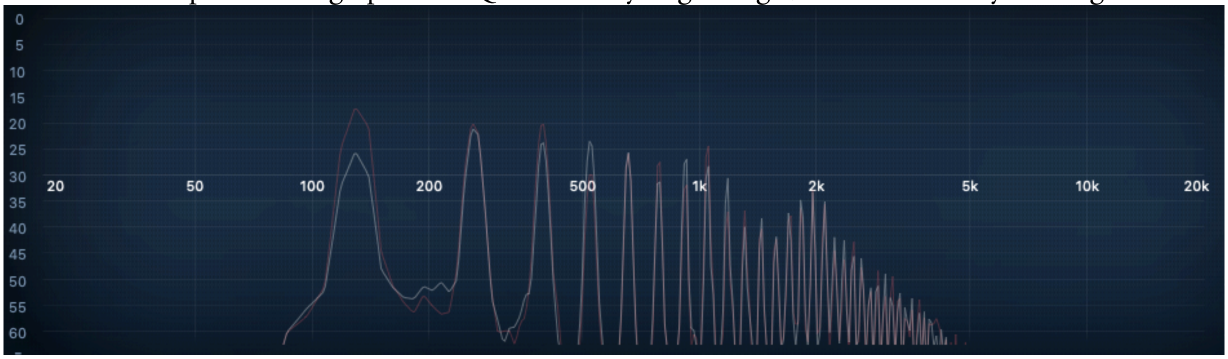

You can clearly see spikes at the fundamental pitch at around 130 Hz, as well as the second harmonic at around 260 Hz (2:1), the third harmonic at around 390 Hz (3:1) and so on, with harmonics generally tapering in amplitude (i.e. volume) as they get higher. You can also see them getting closer together because their ratios are getting smaller, even though their linear distance is even throughout. (There is a 130 Hz difference between each adjacent harmonic.)

Compare that image to a frequency analysis of a sustained "noise," i.e. a sound that we do not perceive as a musical tone:

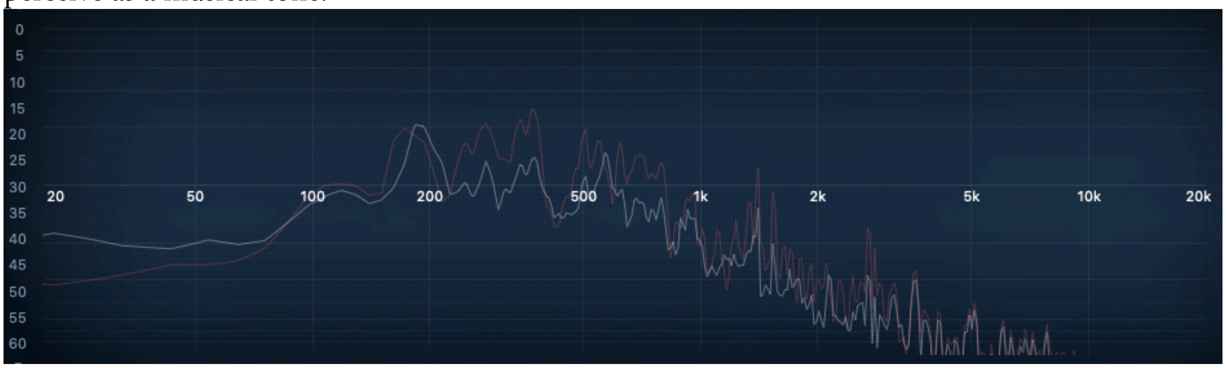

There is no clear pattern of whole-number ratios above the fundamental frequency. The sound is complex and contains many partials, but they do not correspond to the harmonic series, and therefore this sound is "noise," rather than a "note," or **inharmonic**.

Although the overtone frequency relationships are the same for every musical pitch, the relative amplitude of each harmonic is not. In fact, the amplitude balance of the harmonics is what gives us the difference between sounds that are different in "timbre" or "color."

If you compare the sound of a violin sustaining a C_4 to the sound of a clarinet sustaining a C_4, they will sound completely different, even though they are playing the same note. The reason for this is that when either instrument plays, the fundamental frequency (C_4) is supplemented with harmonic overtones (C_5, G_5, C_6, E_6, G_6), but the amplitude (volume) of each of those overtones is different in one instrument than it is in another.

You can simulate these timbre changes and simultaneously irritate anyone within hearing range by sustaining a single musical note with your voice while moving your lips and tongue to change the vowel sound that you are singing. ("aaaaaaaeeeeeeeeoooooooooooiiiiiiiiiiuuuuuuuu…") Better yet, do this into a microphone and turn on your EQ analyzer. If you are sustaining a single note, what you will see is a harmonic pattern like in the image on the previous page, but each of the peaks will change in depending on what vowel you are using. This is because even though you are not changing the pitch (the fundamental frequency), changes in the positioning of your lips and tongue alter which harmonics are being emphasizing by the resonant space.

Understanding the harmonic series is important to composing, especially when it comes to orchestration and electronic music production. When reviewing this lesson, consider the questions on the following page.

Discussion Questions

- If every perceived musical note is actually a collection of pitches with different balances, what does that say about how we perceive two different instruments playing in unison? What about two different instruments playing the same thing an octave apart?

- What role could inharmonic sound play in music?

- What is “distortion” in terms of harmonic content, and how do we perceive it?

- How might you use a multi-band EQ to alter a “noisy” sound to be more “musical”?

- Is a single instrument capable of different “colors,” i.e. can an instrument change the balance of the harmonic overtones on a given note?

- Are you familiar with the instrumental performance technique referred to as “harmonics?” How might this relate to the concept of harmonics discussed in this document?

- (Advanced): If you understand the relationship of harmonics that instruments can play to the harmonic series, are there situations in which the introduction of harmonics in a score might impact the listener’s perception of intonation? In what cases would that be an important consideration, and in what cases would it not?